

1 We claim:

2 1. An apparatus for resource migration, comprising a storage system having  
 3 a plurality of storage servers with a set of resources partitioned thereon, said storage  
 4 servers having a load monitor process capable of communicating with other  
 5 load monitor processes for generating a measure of loading on respective  
 6 ones of the plurality of servers;  
 7 a resource migration process for transferring a resource from one of said plurality of  
 8 servers to another of said plurality of servers in response to said measure of  
 9 loading.

10 2. The apparatus of Claim 1, wherein said servers are equivalent to each other.

11 3. The apparatus of Claim 1, wherein said resources are selected from the group  
 12 consisting of data blocks, program files, multimedia files, applications, and database  
 13 files.

14 4. The apparatus of Claim 1, wherein said measure of loading reflects both a  
 15 storage system load and a server load.

16 5. The apparatus of Claim 1, wherein said storage system is a Storage Area  
 17 Network.

18  
 19 6. The apparatus of Claim 1, wherein the load monitor includes a process to  
 20 determine whether a server is servicing a disproportionate share of the client requests  
 21 being handled by a server group.

22  
 23 7. The apparatus of Claim 1, wherein the resource migration process includes a  
 24 block data migration process.

25  
 26 8. The apparatus of Claim 1, further including a routing table for tracking  
 27 resources maintained on the system.

9. The apparatus of Claim 1, wherein a pointer to a resource is maintained during a an access operation to provide continuous data access.

10. The apparatus of Claim 1, wherein the load monitoring process monitors one or more of network traffic load, I/O request load, storage traffic pattern type.

11. The apparatus of Claim 1, wherein the resource migration process includes a further process to detect when a resource write request applies to a resource that is in the process of being moved from a first server to a second server, and apply such resource write request to both copies of the resource held at said first and said second server.

12. The apparatus of Claim 1, wherein the resource migration process divides the resource being moved into smaller subresources, such that each subresource is moved from a first server to a second server in turn, and recovery from failure requires only the recovery of the subresource being moved at the time of failure and subsequent subresources.

13. The apparatus of Claim 12, wherein the resource migration process includes a further process to detect when a resource write request applies to a subresource that is in the process of being moved from a first server to a second server, and apply such resource write request to both copies of the resource held at said first and said second server.

14. A process for moving resources across a storage system having a plurality of storage servers with a set of resources partitioned thereon, comprising the steps of  
monitoring a load on a server and communicating with other load monitor processes for generating a measure of loading on respective ones of the plurality of servers; and  
transferring, as a function of the measured loads, a resource from one of said plurality of servers to another of said plurality of servers in response to said measure of loading.

15. The process of Claim 14, wherein said servers are equivalent to each other.

16. The process of Claim 14, measuring a load includes measuring a storage system load and a server load.

17. The process of Claim 14, including the further step determining whether a server is servicing a disproportionate share of the client requests being handled by a server group.

18. The process of Claim 14, wherein the resource migration process includes a block data migration process.

19. The process of Claim 14, further including maintaining a routing table for tracking resources maintained on the system.

20. The process of Claim 14, wherein the load monitoring process monitors one or more of network traffic load, I/O request load, storage traffic pattern type.

21. The process of Claim 14, further including maintaining a pointer to a resource is maintained during an access operation to provide continuous data access.

22. The process of Claim 14, further including detecting when a resource write request applies to a resource that is in the process of being moved from a first server to a second server, and applying such resource write request to both copies of the resource held at said first and said second server.

23. The process of Claim 14, further including dividing the resource being moved into smaller subresources, such that each subresource is moved from a first server to a second server in turn, and recovery from failure requires only the recovery of the subresource being moved at the time of failure and subsequent subresources.

24. The process of Claim 23, further including detecting when a resource write request applies to a subresource that is in the process of being moved from a first server to a second server, and apply such resource write request to both copies of the resource held at said first and said second server.

- 1 server to a second server, and apply such resource write request to both copies of the
- 2 resource held at said first and said second server.